Resource Modeling for Multiple Storage Devices in Balanced or Substantially-Balanced Workload Environments

Resource Model Equation (4) or similar relationships between I/O capacity and available buffer space may be implemented as described above in single storage device environments. Such relationships may also be extended to apply to multiple storage disk environments by, for example, factoring in the performance impact of multiple storage device environments into the estimation of average access AA and the average transfer rate TR in situations where the total set of viewers may be thought of as being served sequentially. Alternative methodology may be desirable where information (*e.g.*, content such as a popular movie) is replicated across several storage devices (*e.g.*, disk drives) so that client demand for the information may be balanced across the replicas. In such cases, a resource model may be developed that considers additional system I/O performance characteristics such as explicit parallelism and its performance improvement in the system, in addition to the previously described system I/O performance characteristics such as average access and transfer rate.

In those cases where I/O workload is substantially balanced (*e.g.*, substantially evenly distributed across multiple storage devices or groups of storage devices) or is near-balanced (*e.g.*, where maximum Skew value for any given storage device or group of storage device is less than about 2, alternatively from about 1 to less than about 2), an analytical-based Resource Model approach may be employed. This may be the case, for example, where information placement (*e.g.*, movie file placement) on multiple storage devices (*e.g.*, multiple disk drives or multiple disk drive groups) is well managed. Such an analytical-based Resource Model may function by estimating or otherwise modeling how workload is distributed across the multiple storage devices or groups of storage devices using, for example, one or more system I/O performance characteristics that is reflective or representative of workload distribution across the devices or groups of devices, *e.g.*, that reflects the unevenness of I/O demand distribution among multiple storage devices. In one embodiment, the constant value "Skew" may be employed. As previously described in relation to FIG. 1, Skew reflects maximum anticipated retrieval demand allocation for a given storage device 110 in terms of an even retrieval demand distribution among the total number of storage devices 110.

One resource model embodiment that may be implemented in substantially balanced parallel-functioning multiple storage device environments may employ substantially the same buffer space constraints as described above in relation to single storage device environments. However, such a resource model may be implemented in a manner that considers additional system I/O performance characteristics to reflect the substantially balanced parallel nature of workload distribution in the determination of I/O capacity constraints. These additional performance characteristics may include, but are not limited to, the number of storage devices NoD against which I/O demands are distributed, and the Skew value. For example, if there are total of NoV viewers and there are total of NoD storage devices or groups of storage devices against which the I/O workload is distributed in parallel, then the number of viewers NoV_1 that each storage device or group of storage device is expected to support may be expressed as follows:

$$\boldsymbol{NoV_1 = Skew * (\ NoV/NoD)} \qquad (5)$$

Similarly, the Skew value may be used to estimate or approximate the maximal aggregated consumption rates, ("MaxAggRate_perDevice" or "Max$\{\Sigma_{i\,\varepsilon\mathrm{Device}} \boldsymbol{P_i}$; for all storage devices/groups$\}$") of viewers that are served by each storage device or group of storage devices as follows:

$$\mathrm{Max}\{\Sigma_{i\,\varepsilon\mathrm{Device}} \boldsymbol{P_i};\ \text{for all storage devices/groups}\} \sim \boldsymbol{Skew} * (\Sigma_{i=1}^{Nov} \boldsymbol{P_i})\,/\,\boldsymbol{NoD} \qquad (6)$$

Similar to single storage device embodiments, in multiple storage device embodiments cycle time T should be greater than or equal to the maximal aggregate storage device service time for continuous playback, *i.e.*, the maximal aggregate sum of access time and data transfer time for all storage devices or groups of storage devices. Using the above relationships to balance sufficient I/O capacity and sufficient total available buffer space for multiple storage device environments in a manner similar to that employed for Resource Model Equation (4) for

single storage device environments, range of cycle time T to ensure continuous playback may be defined in one embodiment for substantially balanced multiple storage device environments by Resource Model Equation (7) as follows:

$$(Skew / NoD) * NoV * AA / [1 - (Skew / NoD) * (\Sigma_{i=1}^{Nov} P_i) / TR] \leq T$$

$$\leq B_{max} / (\Sigma_{i=1}^{Nov} P_i) \qquad (7)$$

For an embodiment employing a buffer-sharing scheme, the following equation may be alternatively employed:

$$\boldsymbol{(Skew / NoD) * NoV * AA / [1 - (Skew / NoD) * (\Sigma_{i=1}^{Nov} P_i) / TR] \leq T}$$

$$\boldsymbol{\leq B_{max} / [(1 - B\_Save) * (\Sigma_{i=1}^{Nov} P_i)]} \qquad (7')$$

In the practice of the disclosed methods and systems, Resource Model Equation (7) may be employed for I/O admission control and the read-ahead estimation in a manner similar to Resource Model Equation (4). As with Resource Model Equation (4), the number of existing viewers and their estimated playback rates may be tracked and utilized in system 100 of FIG. 1 by I/O manager 140 of storage management processing engine 105 to determine whether or not system 100 can support all viewers without compromising quality of playback (*e.g.,* video playback). As with Resource Equation (4), if values of I/O capacity and buffer space determined from Resource Model Equation (7) overlap then system 100 cannot support all viewers without compromising quality of playback. However, if a value or range of values for cycle time T exist that will satisfy Resource Model Equation (7), then system 100 can support all viewers. Assuming the latter to be the case, Resource Model Equation (7) may be used to determine a range of cycle time T suitable for continuous playback and to give an estimation of optimal read-ahead size for each viewer in a manner similar to that described for Resource Model Equation (4). Thus, Resource Model Equation (7) may be employed to adjust cycle time T and read-ahead size for existing and new viewers in order to maximize the number of viewers supported by an information management system having multiple storage devices or groups of storage devices.